

09/547223

**Method, System and Program Product for
Managing Network Performance**

Field of the Invention

The present invention relates to management of network services and more particularly to a method, system and program product for managing network performance by supporting generation of reliable, anticipatory alerts of potential performance violations.

Background of the Invention

When any computer network is put into service, the network operator and the network users have their own expectations as to the level of performance to be provided by the network. Where the network operator and the network users work for the same organization, the expectations may be formalized in written memoranda or may exist only in the minds of the network users and (hopefully) the network operator.

Where the network operator and the network users work for different organizations, the expectations may be formalized in a service level agreement. A service level agreement or SLA is an agreement or contract between a service provider, the network operator, and a customer, the network user. Under a service level agreement, the customer pays a service fee in return for an assurance that it will receive network service that conforms to requirements defined by the service level agreement. If the service provider then fails to provide the agreed-to service, it ordinarily becomes subject to penalties under the agreement, such as being required to rebate at least some previously received service fees or being required to reduce fees due for future services.

1 While an almost infinite variety of service level agreements, both technical and
2 non-technical in nature, are possible, the present invention generally relates to the
3 management of network performance where performance requirements have been
4 defined, either informally or in formal service level agreements.

5 Network performance requirements, whether formal or informal, should reflect
6 the type of network service being provided and the customer's specific requirements when
7 it uses that service. A customer with high reliability requirements may, for example,
8 expect or even obligate the service provider to keep the network in operation for no less
9 than a specified percentage of time. Similarly, a customer for whom network response
10 time is critical may expect or obligate the service provider to maintain average network
11 transit times on critical routes at or below a defined threshold.

12 To verify that transit time requirements are being met, the service provider can
13 regularly have a source network station "ping" (query) a destination network station to
14 determine round trip transit time; that is, how long it takes for the query to reach the
15 destination and for an acknowledgment to be returned from the destination to the source.

16 The actual performance of the system is usually monitored by a network
17 management application which generates a message or alert when a performance
18 violation occurs. That alert is sent at least to the service provider to enable the service
19 provider to take steps to restore conforming network operation. This approach, while
20 common, has significant drawbacks for both the network user and the service provider.
21 From the network user's perspective, the performance violation may have already caused
22 disruptions of significant tasks or processes by the time the network user first learns of it.
23 Even if the service provider responds promptly to a violation alert, the recovery time or
24 time required to return to conforming network operation is necessarily prolonged since
25 the service provider can't begin to fix a problem until the problem is known to exist.
26 From the service provider's perspective, the service provider may already be subject to

1 penalties under an existing service level agreement by the time it first learns of the
2 penalty-inducing violation. Even where no formal service level agreement exists, the
3 service provider can expect to lose customer good will for having failed to live up to the
4 customer's expectations.

5 Summary of the Invention

6 The present invention may be implemented as a method, system or program
7 product which supports the reliable prediction of network performance violations so that
8 a service provider receives advance warning of an impending violation and can take steps
9 to avoid the predicted violation.

10 The invention can be implemented as a computer-implemented method of
11 managing network performance where performance requirements have been established.
12 The provided service is monitored on a recurring basis to obtain samples of actual values
13 of a performance-defining metric. A trend in actual service is established based upon the
14 obtained samples. Once the trend is established, the time at which the provided service
15 will cease to meet the established performance requirements if the trend continues can be
16 determined.

17 Brief Description of the Drawings

18 While the specification concludes with claims particularly pointing out and
19 distinctly claiming that which is regarded as the present invention, details of a preferred
20 embodiment of the invention may be more readily ascertained from the following detailed
21 description when read in conjunction with the accompanying drawings wherein:

22 Figure 1 is a schematic representation of a network environment in which the present
23 invention may be implemented;

1 Figure 2 is a block diagram of essential components of a network management station in
2 which the invention may be performed;

3 Figure 3 is a functional flow diagram depicting major operations which take place when
4 the invention is used;

5 Figure 4 is a plot of performance metrics over several sampling intervals;

6 Figure 5, consisting of Figures 5a and 5b, taken together, is a flowchart of essential steps
7 performed by a method implementing the present invention;

8 Figure 6 is a plot of conditions under which a pending alert can be canceled for certain
9 successive network performance trends;

10 Figure 7 is a plot of conditions under which a pending alert can be canceled according to
11 an alternate embodiment of the invention; and

12 Figure 8 is a partial flow chart showing method steps that are performed in implementing
13 the alternate embodiment of the invention.

14 Detailed Description

15 Referring to Figure 1, the present invention is used in the administration of
16 computer networks, one example of which is a network 10. The network 10 is
17 represented as including a wide area network 12 which connects local networks to
18 remote networks (not shown). The interface between the local networks and the wide
19 area network 12 is provided through a gateway device 14 having an attached network
20 management workstation 16. The illustrated local networks include both a token ring
21 local area network (LAN) 18 and an ethernet LAN 26. Token ring LAN is shown as

1 having network stations 20 and 22 and a bridge 24 to the gateway device 14. Ethernet
2 LAN 26 is shown as including network stations 28 and 30 and a bridge 32 to the gateway
3 device 14.

4 The types of networks and network devices shown in the drawing are intended as
5 examples of a suitable environment for the present invention. The invention can be used
6 in virtually any multi-node network where a performance metric is measurable. The
7 invention should in no way be considered to be limited to the illustrated environment.

8 Specific embodiments of the invention will be described below, but it should be
9 kept in mind that the present invention can be implemented in several different forms,
10 such as in special purpose hardware or in a combination of hardware and software. A
11 typical combination of hardware and software is a general-purpose computer system
12 using a computer program that, once loaded and executed, causes the system to carry out
13 method steps which will be described below. The software may be pre-loaded into the
14 general-purpose computer system or may be separately available as a computer program
15 product which, when loaded into a computer system, causes the system to carry out the
16 methods steps.

17 The term "computer program" in the present context means any expression, in any
18 language, code, or notation, of a set of instructions intended to cause a system having
19 information processing capability to perform a particular function either directly or after
20 conversion to another language and/or reproduction in a different material form.

21 Figure 2 illustrates the major physical components of a general-purpose computer
22 system capable, when programmed properly, of implementing the present invention. The
23 computer system includes a central processing unit (CPU) subsystem 34 with a processor
24 and supporting registers, caches and logic circuits. The computer system further includes
25 random access memory 36, hard drive 38 and an optical drive 40, such as a CD/R,

1 CD/RW or DVD drive. Where the invention is implemented as a program product, it is
2 typically made available to the network operator initially on removeable magnetic or
3 optical media for installation onto hard drive 38. Once the initial installation is complete,
4 the program can be transferred into random access memory 36 as needed from hard drive
5 38. Alternatively, the program may be loaded into random access memory 36 directly
6 from an optical media mounted in optical drive 40. The computer system further
7 includes system input/output (I/O) adapters 42 supporting connections to standard system
8 components such as a keyboard 44, a pointing device 46 and a display monitor 48.
9 Finally, the computer system includes a network interface card 50 which provides the
10 needed interface to the rest of the network.

11 Figure 3 is a functional flow diagram illustrating major functions that are
12 performed by a computer system programmed in accordance with the present invention.
13 Among other tasks, the computer system functions as a network performance monitor 52
14 by making and/or receiving measurements reflecting actual network performance over
15 time. The performance measurements constitute samples which are processed by a
16 service metric sample processor function 54 to convert those samples to a metric (such as
17 an average value) which reflects current network performance. For the sake of simplicity,
18 the following discussion assumes that a single type of metric (average ping time) is
19 monitored. In some situations, it may be desirable to monitor more than one type metric
20 so that appropriate actions can be taken where any one of the metrics exceeds an
21 allowable value.

22 Where successive values for a defined metric have been gathered, those values
23 can be processed in a trend module generator to determine whether there is a recognizable
24 trend in the metric values over time. Where a metric is trending toward an unacceptable
25 value, an alert generator function 58 can generate and send an anticipatory alert to the
26 service provider in advance of an actual violation. The anticipatory alert gives the service

1 provider time to take steps which will head off an actual violation of defined performance
2 requirements.

3 Figure 4 is plot of a specific service metric over several sampling intervals. The
4 specific service metric is ping time on a particular route between a first network station
5 and a second network station; i.e., the time required for the first station to send a ping or
6 query to the second station and to receive a response from the second station. Typically,
7 the first station, which may be a network management station, is required to perform a
8 minimum number of ping tests over a standard sampling interval 60 which, for purposes
9 of this description, is assumed to be a 24-hour day. The actual or raw samples gathered
10 over the course of each sampling interval can be processed to obtain an average ping
11 value representing the average network performance over the entire day. Object 62
12 represents the average ping value over a first sampling interval 60. As a matter of
13 convention, object 62 shown as occurring at the midpoint of the interval even though its
14 value can't be determined until the interval has ended. To establish a trend in actual
15 network performance, ping times are taken throughout the day and are averaged to
16 establish the actual network performance for that day. Objects 64 and 66 represent the
17 ping time averages for the second and third sampling intervals on the plot.

18 A trend in actual network performance is established by using two or more of the
19 average ping time values and known linear regression techniques to derive a curve or line
20 68 representing the trend. Depending upon the service metric chosen and the network
21 performance requirements, the trend-indicating line can be a simple straight line
22 established using two acceptable metric averages or a curved line fitted using several
23 successive acceptable metric averages. Assuming a straight line 68 adequately describes
24 the trend, the slope (positive or negative) of that line indicates whether the actual network
25 performance over time is trending toward or away from a limit 70 of acceptable network
26 performance (maximum allowable average ping time).

1 A simple approach to network management would be to wait for the average ping
2 time to exceed limit 70 before generating any sort of alert to the system provider. The
3 present invention uses a better approach. The slope of the trend-indicating line can be
4 calculated using two acceptable ping time averages. Once the slope of the trend-
5 indicating line and at least one average ping time value is known, straight-forward
6 mathematic calculations can be used to predict the time t_v at which the average ping time
7 will exceed the limit 70 if the trend continues unchanged.

8 In accordance with a preferred embodiment, an alert is not sent simply because a
9 trend toward unacceptable ping times is established. For an alert to be of interest to a
10 system provider, it must be reasonably imminent. A system provider is not likely to want
11 to respond to a prediction of unacceptable ping times far in the future given the possibility
12 that the trend toward unacceptable ping times might level off or be reversed in the course
13 of normal system operation. For that reason, an alert is generated and sent to the system
14 provider only where the predicted violation time t_v falls within a time window (for
15 example, two days) beginning at the current time. If t_v is predicted as occurring outside
16 of the time window, no alert is generated.

17 Even where an alert has been generated and sent to the service provider, the
18 possibility still exists that the trend toward increasing ping time averages will level off or
19 reverse itself in the course of normal system operation. In accordance with one feature of
20 the invention, network performance (represented by average ping time) continues to be
21 monitored even after an alert is generated. If the trend resulting in a pending alert is
22 found to have changed substantially, a pending alert may be canceled.

23 For the described process to work reliably, the data used in the process must be
24 reliable. In any process which relies on sampling of actual values, there is always the
25 possibility that abnormal system conditions will result in abnormal sample values during
26 any given sampling interval. To eliminate unreliable sets of samples, the present

invention imposes reliability tests for each set of samples used in establishing a performance trend. If the reliability tests, described below, are not satisfied for a particular set of samples, the set is ignored, at least for trend determination purposes. The set of samples may be retained in the system for other purposes beyond the scope of this invention.

A first and seminal reliability test is that the number n of samples obtained over a sampling interval must exceed a predetermined minimum. Conventionally, it is assumed that at least thirty measurements or samples of a particular metric are needed to support reliable statistical analyses. If, during a particular sampling interval, less than thirty samples are obtained, no attempt is made to establish a performance trend using the sample set.

Assuming the necessary minimum number of samples have been obtained over the sampling interval, a second reliability test uses standard statistical techniques to derive the statistical mean and the statistical standard deviation of the set of samples under consideration. For a set of n samples, each having an individual raw value y_i , the statistical mean is simply the average of the values; that is

$$y_{mean} = \frac{\sum y_i}{n}$$

For the same set of samples, the standard deviation s can be computed as

$$s = \sqrt{\frac{n \sum (y_i)^2 - (\sum y_i)^2}{n(n-1)}}$$

1 To determine whether a particular set of samples, the mean value y_{mean} and the
2 standard deviation s of the set are used to generate a Confidence Percentage value CP
3 where

$$4 \quad CP = \frac{y_{mean}}{s}$$

5 A set of samples is considered reliable (and thus suitable for use in the described process)
6 if CP does not exceed a predetermined percentage threshold, preferably on the order of
7 25%. If CP exceeds the predetermined threshold, no effort is made to determine a
8 performance trend based on the "unreliable" set of samples.

9 Where a set of samples gathered during a particular sampling interval are not to be
10 used as failing to meet reliability tests, acceptable samples gathered during preceding and
11 following sampling intervals can still be used to establish the trend in network
12 performance.

13 Figure 5 is a flowchart of method steps that are performed in implementing the
14 present invention. The initial step 74 is to perform a system test which generates raw
15 sample values. The ping response time test described above is just one example of many
16 types of system tests which might be performed to obtain a measure of actual network
17 performance. Each test 74 is followed by a time check 76 which determines whether the
18 current sampling interval has just ended or expired. If the sampling interval has not
19 expired, a second time check 77 is made to determine whether a inter-sample interval
20 timer has expired.

21 The inter-sample interval timer is used to limit the number of samples acquired
22 during a given sampling interval since every test operation performed to acquire a sample
23 represents network overhead and necessarily impacts network throughput. As noted
24 earlier, good statistical practice requires a minimum of thirty samples for statistically

1 reliable averaging. Therefore, it can be expected that the inter-sample interval timer will
2 have a short enough timeout period to guarantee that at least thirty samples will be
3 obtained over the course of the sampling interval. The maximum number of samples to
4 be obtained may vary with the type of system test being performed. For ping time tests,
5 it is believed that a maximum of 130 - 150 samples per twentyfour hour sampling interval
6 is appropriate.

7 When the sampling interval does expire, the number of samples obtained during
8 the interval is compared to a minimum threshold number in operation 78. If the number
9 of samples falls below the minimum threshold, no effort is made to continue the trend
10 determination process and the current process cycle is ended. Even though the current
11 process cycle ends, a new sampling cycle is already underway for the new sampling
12 interval that has just begun.

13 Assuming an adequate number of samples is obtained for the current cycle, the
14 raw samples are summed in step 80. In a following step 82, each raw sample in the set is
15 squared and the squared values are summed. The average or mean value for the set is
16 obtained in step 84 while the standard deviation for the set is calculated in step 86.

17 As described earlier, the set of samples may or may not be used depending the
18 confidence percentage *CP* for the set; that is, the ratio of the set's standard deviation to
19 its mean or average value. The *CP* value is calculated in step 88 using the earlier-
20 described equation and then compared to a predetermined threshold percentage in step 90
21 to determine whether the set's *CP* value falls within acceptable limits. If the set's *CP*
22 value falls outside the acceptable limits, the trend determination process is ended without
23 using the "unreliable" set of samples

24 Assuming the set of samples under consideration satisfies the defined reliability
25 tests, the averages or mean values for the current set and an earlier set of samples are used

1 in an operation 94 to determine whether there is a trend in average ping response times.
2 The trend is characterized by the slope of a line passing through the two time displaced
3 mean values. The slope is tested in step 96 to determine whether the average ping
4 response times are approaching a violation threshold. If step 96 shows that the trend is
5 toward violation, the current slope of the line, one of the average ping response times at
6 an endpoint of the line and the violation threshold are used to predict (step 98) when the
7 average ping response time will exceed the threshold assuming the current trend
8 continues unchanged.

9 This predicted time-until-violation value can be determined by solving the
10 equation

11 $y = mx + b$ for the value of x where

12 y = the maximum acceptable (violation threshold) average ping time,

13 m = the computed slope of the trend line during the last sampling interval,

14 b = the current average ping time, and

15 x = the time-until-violation as measured from the current time.

16 The variables y , m and b are known, making it a simple matter to determine x . Once the
17 predicted violation time is established, it can be checked in an operation 100 against the
18 limits of a time window (for example, a time window that begins at the current time and
19 ends 48 hours later). If the predicted time of violation falls outside the time window, the
20 current process cycle is ended with no action being taken other than to preserve the
21 values calculated using the current set of samples. However, if the predicted time of
22 violation falls within the time window, an alert is generated in step 102 and sent to the
23 network manager.

24 If step 96 does not indicate that the current trend is toward the violation threshold,
25 meaning the trend is either flat or away from the violation threshold, then a check 103 is
26 made as to whether a previously generated alert is still pending. If there is no pending
27 alert, no further computations are performed and the current process cycle is ended.

1 If a previously generated alert is still pending, the absolute value of the slope of
2 the current trend line is compared to the absolute value of the slope of the preceding trend
3 line in an operation 104. Unless the absolute value of the new slope is greater than the
4 absolute value of the preceding slope while the sign of the new slope is negative, the
5 trend toward an eventual violation necessarily continues. The samples and the metric
6 average are retained. The previously-generated alert is not affected. The current process
7 cycle is ended to allow the next iteration of the process to continue.

8 If, however, the absolute value of the new slope is greater than the absolute value
9 of the old while their algebraic signs are different, a significant trend away from the
10 violation threshold is necessarily indicated. . This can most clearly be seen by reference
11 to Figure 6 where line 110 represents an old or prior trend line while line 112 represents
12 the current trend line. While the slope of line 110 shows a trend toward violation, the
13 slope of line 112 shows an even sharper trend away from violation. Referring back to
14 Figure 5, where a significant trend away from violation is found from the test 104, the
15 previously-generated and still pending alert is canceled in step 106. }

16 An alternative and less stringent test for determining when to cancel a previously
17 generated alert is described below with reference to Figures 7 and 8. The alternative test
18 is based on a premise that a pending alert issued as a result of a prior trend can safely be
19 canceled if an alert would not be generated based on the current trend. Recall that an alert
20 is generated in the process described above where a trend toward a violation threshold
21 will cross that threshold within a predetermined time window if the trend continues
22 unchanged. A time window of two days was assumed for purposes of illustration.

23 Referring first to Figure 7, which illustrates the premise of the alternative process,
24 an alert is generated at time t3 because the trend characterized by line 114 would result in
25 the violation threshold being exceeded within two days of time t3. However, for the
26 current trend represented by line 118 (beginning at time t3 and ending at time t4), it can

1 be seen that the lesser slope of the current trend would not, if continued, cause the trend
2 to reach the violation threshold will not be exceeded within two days of time t4 even if
3 the trend continues unchanged. Under the noted conditions, no alert would be issued at
4 time t4. If an alert would not be issued at time t4 based on the then current trend, it
5 would be illogical to allow a previously-generated alert to remain in force. If a
6 determination is made that current conditions do not warrant generation of an alert at
7 current time t4, then pending alerts based on past conditions are canceled.

8 Figure 8 is a flow chart of the method steps required to carry out the alternative
9 process steps noted above. The method steps previously described with reference to
10 Figure 5 remain unchanged from the beginning of that Figure through the output from
11 operation 94, which is the slope of the current trend line. In the alternative process, the
12 determined slope is used as an input to a step 120 which determines whether the current
13 trend is toward violation. If it isn't, any pending alerts are canceled. If the trend is found
14 still to be toward violation, the time at which the trend will result in a violation is
15 predicted in step 124. If the predicted time of violation falls is found to fall within the
16 time window in step 126, then a new alert is generated in step 128. Previously-generated
17 alerts (if any) are not canceled.

18 If, however, the test 126 indicates that the latest predicted time of violation falls
19 outside the time window, which means that no alert is to be generated based on current
20 conditions, test 130 looks for previously-generated and still pending alerts. If any such
21 alert or alerts exist, they are canceled in step 132.

1 While there has been described what is believed to be a preferred embodiment of
2 the invention, variations and modifications in the preferred embodiment will occur to
3 those skilled in the art. Therefore, it is intended that the appended claims shall be
4 construed to include the preferred embodiment and all variations and modifications as fall
5 within the true spirit and scope of the invention.

001140-6224560